

Master Thesis Project Proposal (30hp):

Prediction of Task Runtimes and Power Consumption for SkePU Skeleton-Based Parallel Streaming Applications

Many static scheduling approaches, e. g., *Crown Scheduling* (Melot et al. 2015), require information on the runtime behavior of tasks and the power consumption their execution causes for certain degrees of parallelism and processor core operating frequencies. Accurate data can be obtained via microbenchmarking in individual cases (Litzinger and Keller 2022). Unfortunately, it necessitates a large resource investment due to the myriad of combinatorial possibilities, which can generally only be justified for long-running or widely deployed applications. Thus, there is a desire for abstraction, allowing one to describe a task-based application in terms of, e. g., the SkePU skeletons it makes use of as well as the inputs it processes. This information could then be leveraged to predict the aforementioned quantities on a particular system. Evidently, the underlying data will still have to be gathered via microbenchmarking and subsequently refined by means of adequate performance/power modeling, but it will serve to cover a broad range of applications instead of just a single one, significantly easing the deployment of SkePU-based applications when it comes to scheduling and resource management.

SkePU: <https://skepu.github.io/>

N. Melot, C. Kessler, J. Keller, P. Eitschberger (2015): Fast Crown Scheduling Heuristics for Energy-Efficient Mapping and Scaling of Moldable Streaming Tasks on Manycore Systems. *ACM Trans. Archit. Code Optim.* 11:4. <https://doi.org/10.1145/2687653>

S. Litzinger, J. Keller (2022): Code generation for energy-efficient execution of dynamic streaming task graphs on parallel and heterogeneous platforms. *Concurrency Computat Pract Exper.* 34(2):e6072. <https://doi.org/10.1002/cpe.6072>

Task This project aims at generating accurate runtime and power consumption estimations for parallel tasks in SkePU-based streaming applications. Predictions shall be based on high-level task descriptions and a suitable set of microbenchmarks. It will be necessary to identify relevant task properties and to develop a reliable prediction method. Furthermore, the microbenchmarks serving as an empirical basis will have to be curated and carried out on a distributed test system. Ultimately, the estimator could be inserted into the system's resource manager to improve prediction quality.

Prerequisites TDDD56 Multicore and GPU Programming (highly recommended), TDDE65 Programming of parallel computers (recommended), familiarity with C/C++, operating systems (especially Linux), Unix shell (e.g., Bash), ideally experience with SkePU or skeleton-based frameworks in general

Contact Sebastian Litzinger, Christoph Kessler (first.last@liu.se)

Open thesis projects: <https://www.ida.liu.se/~chrke55/exjobb/open-exjobb-projects.shtml>